

Offline

Ray Culbertson
for the Offline Group

Rick StDenis, co-leader
Ray Culbertson, co-leader
Elena Gerchtein, Assoc Head for Production
Stephan Lammel, Assoc. Head for Services
Aidan Robson, Assoc. Head for Projects

Ending
Starting

Production

Elena Gerchtein

Calibration

Willis Sukamoto (L)

Roman Lysak

Dan Goldin

calibrators

Production

Elena Gerchtein (L)
operators

Ntupling

Sasha Golossanov (L)
operators

Monte Carlo

Costas Vellidis (L)
Oksana Tadevosyan
Liqun Li
MC Reps

Infrastructure

Stephan Lammel (L)
(CD)

Code Management

Jim Bellinger (L)
Lynn Garren
(Simona Rolli)
(Rick Snider)
(Donatella Torretta)

Remote Sites

Rick StDenis(L)
Rick Snider (L)
Site Coordinators

Services

Stephan Lammel

Mike Kirby

Grid/DH

Joe Boyd (L,CAF)
Eric Wicklund (L,DH)
Dmitry Onoprienko
Alexei Varganov
(Robert Illingworth)
(Marc Mengel)
(REX)
(Site Coordinators)

Databases

Eric Wicklund (L)
(Barry Blumenfeld)
(Dennis Box)
(DBAs and admins)

Calibrators

SVX Align Dominik Horn
SVX Timo Aaltonen
COT **John Paul Chou**
Kevin Burkett
dE/dx Shin-Shan Eiko Yu
Beamlines Roberto Carosi
TOF Jesus Manuel Vizan Garcia
Barbara Alvarez
PES Halley Brown
PEM Willis Sakumoto,
CHA Fabio Happacher
CEM Larry Nodulman
Cal Timing Adam Aurisano,
CP2/CCR Sam Hewamanage
PASS Tom Riddick
Dan Beecher

Operators

Elisa Pueschel
Olga Terlyga
Oksana Tadevosyan
Zhenbin Wu
Jon Wilson
Aristotle Calamba

MC Reps

HDG **Shalhout Shalhout**
Roman Lysak
TOP Dave Mietlicki
EWK Maria D'Errico
BOT Satyajit Behari
EXO John Strologas
QCD **Erik Jens Brucken**
Manuel Mussini

Site Coordinators

Fermigrid Steve Timm
LCG/CNAF Donatella Lucchesi
Gabriele Compostella
Matteo Bauce
PACAF Masakazu Kurata
Yuji Takeuchi
Suen Hou
Tsan Hsieh
KISTI Seo-Young Noh
Beob Kyun Kim
MIT Max Goncharov
Wisconsin Will Maier
Florida Yu
McGill Adrian Butazu

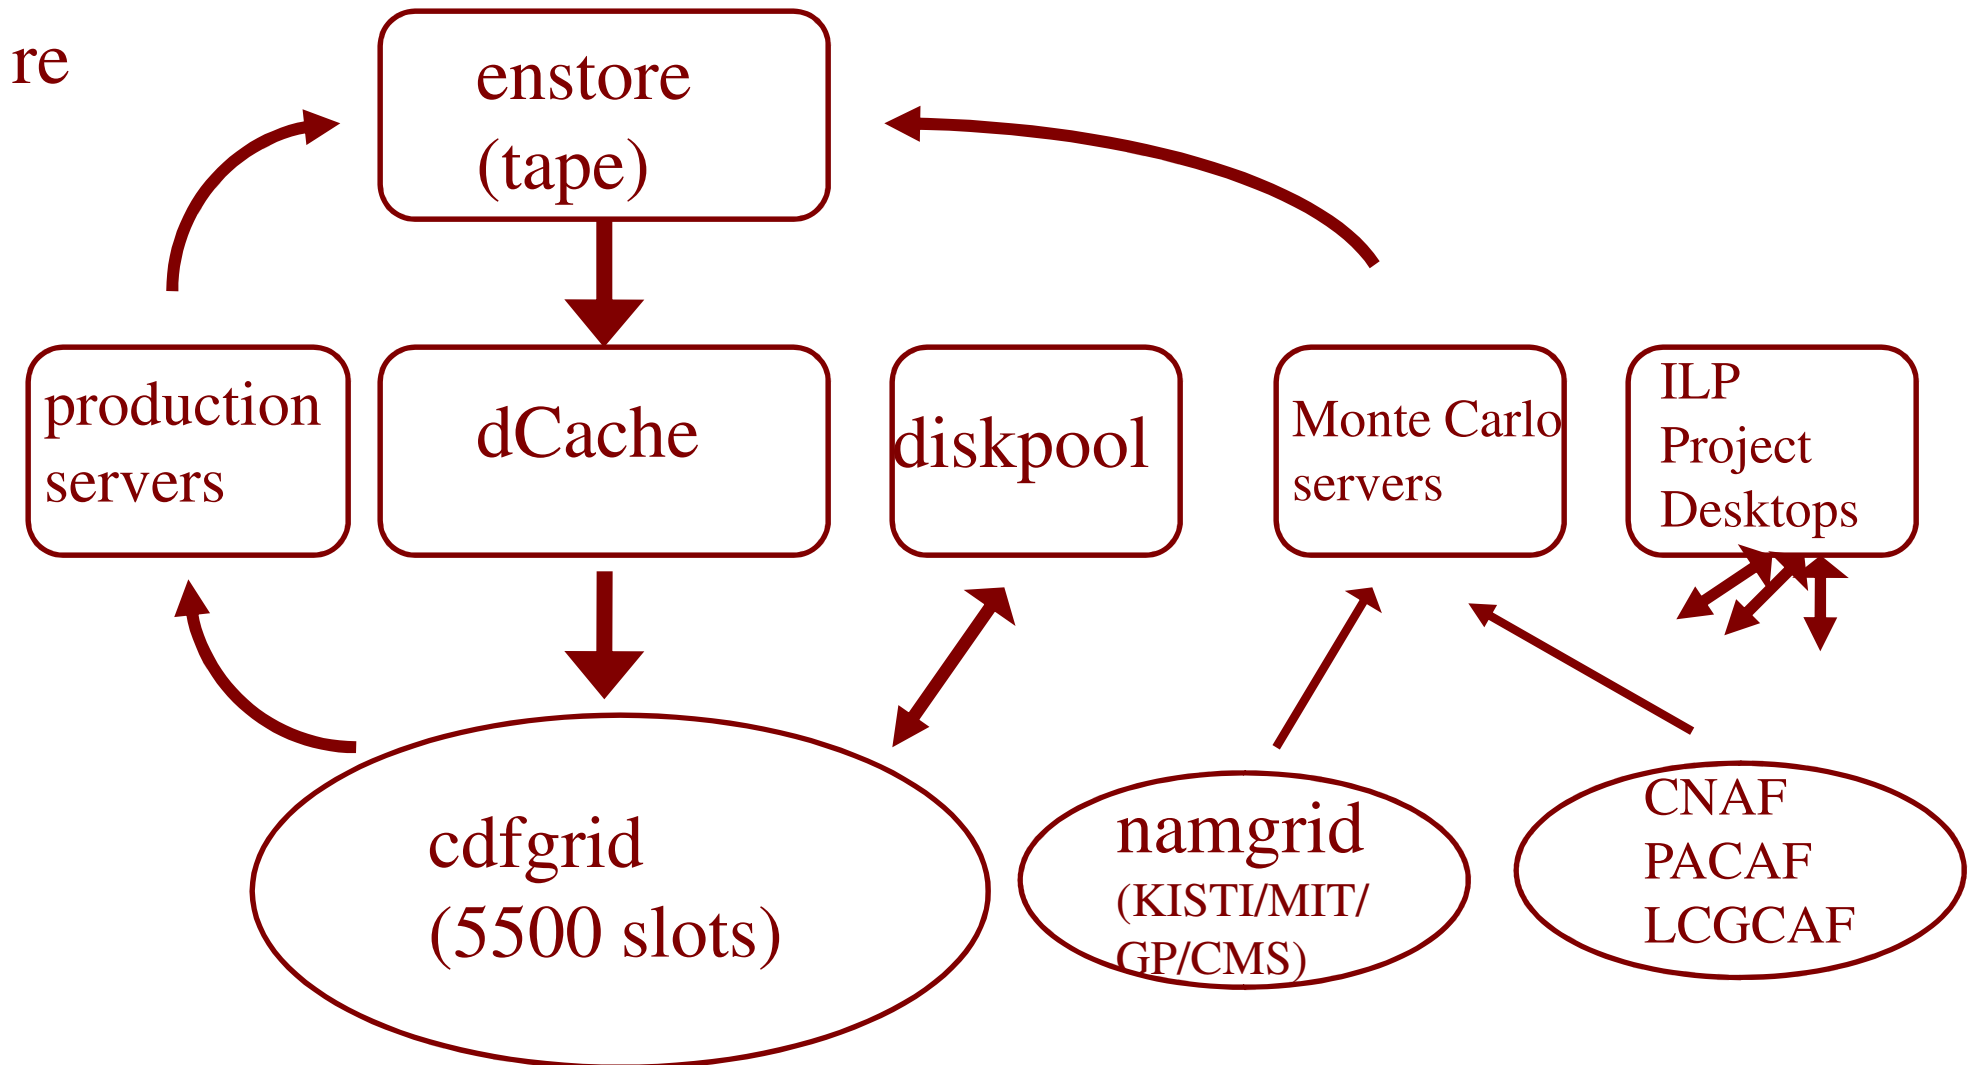
Ending
Starting

SAM Shifters

John Yoh
Ivan Vila
Giovanni Piacentino
Stefano Giagu
Barry Blumenfeld
Teresa Rodrigo
Peter Bussey
Thomas Kuhr
Alberto Ruiz
Aidan Robson




Major Systems Overview





Data Handling

- ▶ Enstore tape system: 16 LTO3 drive, 26 LTO4 drives
 - Delivering typically 20TB per day, 50TB
- ▶ dCache, main disk cache, ~400TB, with tape backend
 - Delivering typically 50TB per day, 150TB peak



- ▶ backlogs in March, OK since then!



Data Handling

► Planning needs

- no new drives added last year
- heavy production load expected next year (see later..)
- new tape storage robot arriving this year, timing a concern
- LTO3→4 migration starting, LTO5 available in Spring



► Concentrate on efficiency

- prestaging by SAM stations - levels load, helps CPU
- one –time reads: ntupling doesn't need files to stay in cache
- coordinating requests by tape
- prestaging re-processing offsite – levels load
- large files (see code management later...)
- dCache will grow by 400TB →800TB in Feb
 - potentially a huge help in reducing tape load!!



CdfGrid

- 5500 slots
 - Smooth operations!
 - small increase in 2010
- Loads:
 - Spring: heavy
 - Summer/Fall: light




< 1 year >



NamGrid

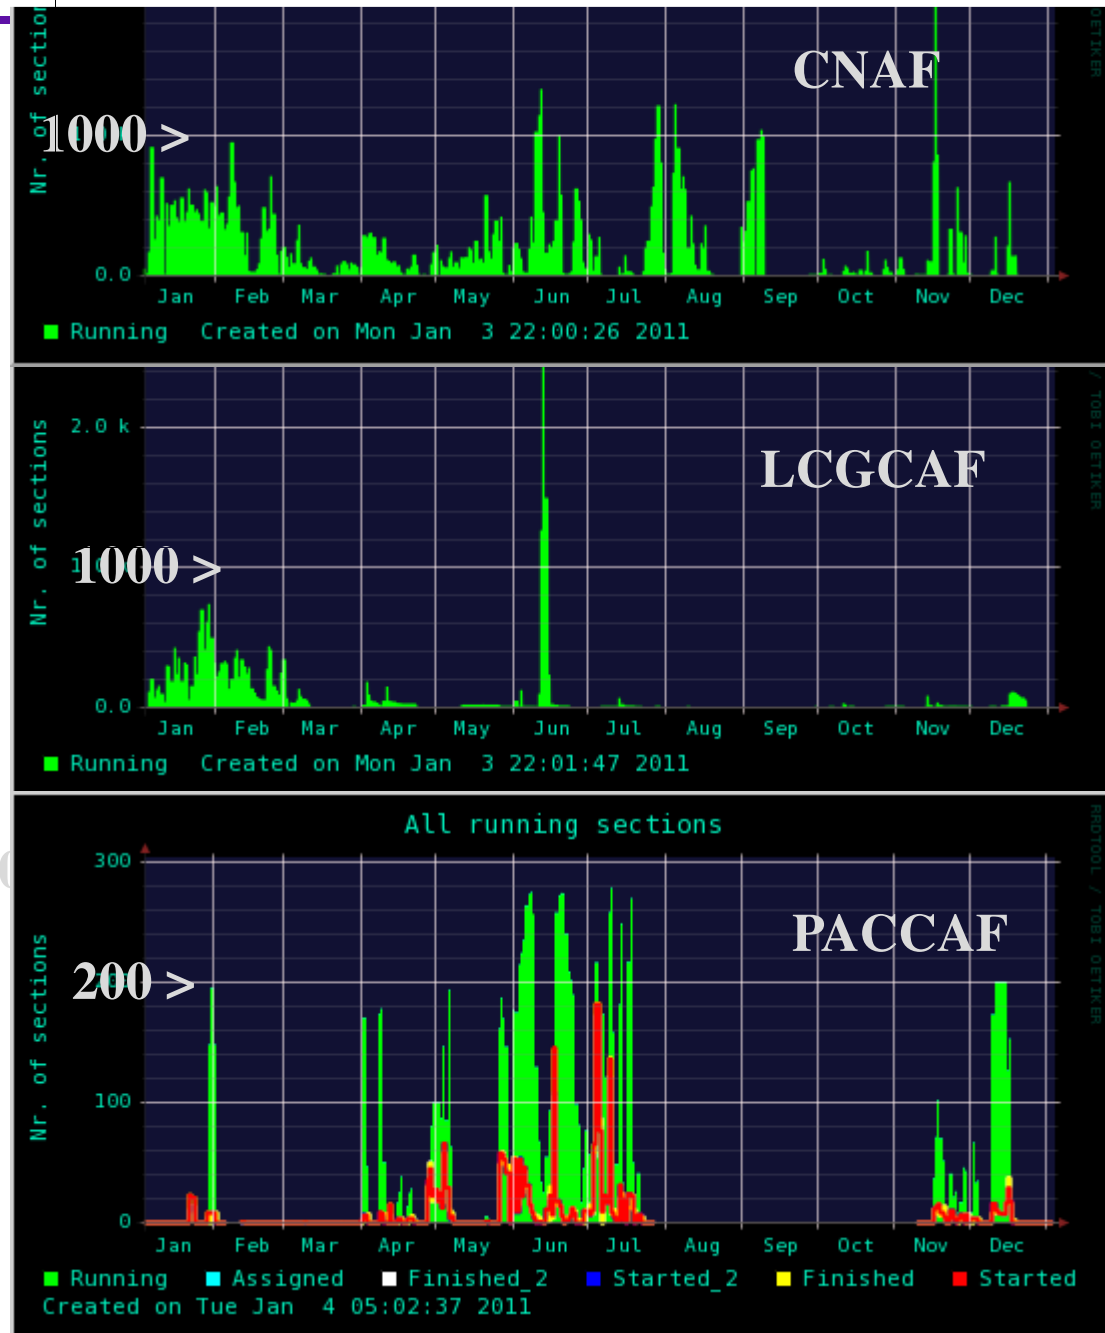
- ▶ A portal to offsite farms running OSG, 15% of CdfGrid
 - achieved regular access to GP and CMS farms!
 - achieved regular backfill to CdfGrid
- ▶ Major achievement – solid integration of KISTI site!
 - accessing SAM cache and CdfCode
 - staging data (200→300TB) and running re-processing
 - regularly runs MC
 - growing..
- ▶ MIT site continues to be very reliable!





Off-site Farms

- ▶ 3 portals
 - direct subm. to farms
 - up to 20% of CdfGrid
- ▶ CNAF
 - reliable, good data transfers
 - possible re-production site!
- ▶ LCGCAF
 - struggles with reliability and data transfers
- use only by experts
- ▶ PACCAF
 - reliable, good data transfers (head node down a while...)





Diskpool

- ▶ 325TB of “persistent” dCache assigned to physics groups
- ▶ mostly smooth operations
 - a few outages for hardware failures
 - forced upgrade in June – very successful!
 - occasional heavy loads
- ▶ Alexei Varganov, our only diskpool expert has left for his next job - good for him, bad for us!
- ▶ Physics group have agreed to try to migrate off diskpool
 - purge stale data
 - little-accessed data copying to archiving tape
 - large datasets – migrating to enstore/dCache
 - smaller or personal datasets migrating to project disk
 - finishing this will be a major project next year!



Production Chain

► Calibration

- 3-4 weeks from end of data period to calibrations ready
- need to reduce one-off delays, and maintain expertise

► Production

- 2-3 weeks to process a period, 1 week to upload/clean-up

► Ntupling

- 2-3 weeks to Stntuple, Topntuple, and BStntuple
smooth operations...

► Production and Ntupling are now being done by operators

► Transitioned to SL5, constantly improving efficiency

► *Production has been successfully run on KISTI and is almost commissioned on CNAF – a major increase in flexibility!*



Monte Carlo Operations

Generated 400M events, 210TB last year!

- that's a lot, but demand was down from previous years!

- ▶ Database overloads
 - lots of work on mitigation and debugging
- ▶ Maxopt
 - commissioned – saves ~50% in CPU
- ▶ cdfSim crashes
 - occur at ~1%, finally being tracked down
- ▶ Absorbing B MC into high-pt tarball – easier maintenance,
B MC gets automatic concatenation
- ▶ SL5 commissioning (see later)
- ▶ major efforts in commissioning and validating farms



Code Management

Huge effort of many people to migrate off SL3 Redhat and commission SL5 - success!

- ▶ SL5 migration required major modifications to offline
 - new infrastructure, new releases
 - discovering issues, fixing and validating cdfcode
 - migration of farms
 - migration of users
 - migration of desktops off SL3
- ▶ Now taking on another large effort: large file support
 - could improve tape speed access by up to $\times 3$
 - can read large files now, need patches to prod. tarballs to write
 - hope to declare validated and deploy very soon!



Reprocessing

The search for the standard model Higgs continues to be a central goal of CDF and Fermilab...

In this search, progress comes from many small improvements.

- ▶ About half of our data has non-optimal Si clustering, reduces tagging 5-10% (Periods 18-28 out of 32)
- ▶ Resurgence in improved tracking and tagging options

Over the last year, it is has started to come into focus how reprocessing data can help.

4 steps to a final Higgs dataset...



Reprocessing

re-process and ntuplize P18-28 to fix Si clustering

1) Phase I

- high priority leptons and jets

- Start Aug 17, prod. done

- ntupling 90% done

- *on target for low mass Higgs update for Summer Conferences!*

2) Phase II - low priority datasets

- Started this Fall at lower priority, 25% done

Production							
Period	b	g	c	j	h	e	
18	c	c	c	c	c	c	Phase 1
19	c	c	c	k	i	c	
20	c	c	k	k	i	c	Phase 2
21	c	c	k			c	
22	k	c	k			c	
23	c	c	k			c	
24	c	c	k			c	
25	c	c	k			c	
26	c	c	k				
27	c	c				c	
28	c	c				c	
29							
30	c	c	c	c	c	c	
31	c	c	c	c	c	c	
32							



Reprocessing

two more steps, more uncertainty in details and schedule, though likely on the year time frame to be useful

- 3) Re-ntuple: P0-17, probably also P18-32
 - capture current tagging improvements
 - significant upgrades to ntuples for tagging, leptons and jets
- 4) Best Ntuple : run on the full dataset
 - if new tracking algorithms can provide fundamental improvements to tagging and vertexing
 - May include parts of production, or not, may save production output, or not, will be at least new ntuples



Two More Projects

Two more projects targeted for B physics...

► BStntuple

- re-ntuple all the B physics data streams
- add covariance matrix for tracks -
 - allows arbitrary vertexing choices at the ntuple level
- enables many new analyses, and is flexible for the future

► Generic B MC

- generate, simulate, produce, ntuple bb Monte Carlo
- has been a wanted for a long, now becoming more urgent
- details and targeted dataset size are under design



Major Projects Summary

	Output (TB)	CPU(1K slotdays)
Phase I	442	50-100
Phase II	728	100-200
Re-ntuple	537	240
Best ntuple	700-2700	240-480
BStntuple	100	60
bb MC	500-1000	20-400
Total	2700-5500	530-1600

- for comparison, typical yearly tape written = 1000
total CPU on CdfGrid per year = 2000

this will be a challenge !



The Bottom Line

- ▶ CDF Offline had a very successful year!
- ▶ Next year could be very challenging if we need to implement all plans
 - might need to take resources from analysis
 - really stretch to use all remote resources, and add more
 - purchase or recycle data tapes
- ▶ Run III
 - current budget, plans are gliding towards end of life
 - can steer back to long-term ops with solid funding

Perhaps even more than usual, IFC contribution is crucial!